Lessons from Concatenated Context-Selected Linear Mappings for Understanding Animal Neural Networks

Introduction

Our experiments and discussions around concatenated context-selected linear mapping systems suggest several constraints that may be relevant well beyond the particular systems we have been studying.

The central idea is that a system can become computationally very complex by concatenating relatively simple transformations, provided that **context selects appropriate transformations** and that those transformations remain sufficiently geometrically compatible.

This leads to an interesting question:

> What might these observations tell us about how animal nervous systems achieve enormous computational complexity without simply degenerating into high-dimensional chaos?

The following is a speculative interpretation rather than a claim that animal brains literally implement concatenated context-selected linear mappings. The value is in the possible constraints and principles that the experiments reveal.

---

1. Meaningful context creates islands of stability

In a context-selected linear mapping system, context is not merely additional information. It determines **which local computation is appropriate**.

If context states are arbitrary, the system has little reason to settle into useful regions. But if context corresponds to persistent environmental or internal states—walking, hunting, threatened, hungry, approaching food, carrying something, nighttime, social interaction—then the system acquires something like **behavioral islands**.

Within such an island:

- similar inputs should produce related responses,
- mappings should remain relatively stable,
- small perturbations should not completely change the computation,
- transitions between islands can be relatively sharp.

This sounds quite similar to what we observe behaviorally in animals.

An animal does not appear to continuously compute an entirely new function from scratch. It often enters a **state of being** in which a large amount of its nervous system behaves differently.

One useful interpretation of neural context is therefore:

> **Context defines a temporarily stable computational regime.**

This makes things such as neuromodulation, recurrent activity, hormones, sensory state, arousal, locomotor state, and memory particularly interesting. They can all provide variables that determine which computational regime the rest of the nervous system operates in.

---

2. Composition strongly favors linear or near-linear transformations

This is perhaps the more interesting lesson.

Suppose we repeatedly compose arbitrary nonlinear transformations:

$$x_{n+1} = f_n(x_n)$$

In a sufficiently high-dimensional space, there is enormous opportunity for the geometry to become distorted. Small differences can be stretched, folded, collapsed, and redirected.

After many stages, the original structure of the signal can become almost unrecognizable.

With approximately linear transformations,

$$x_{n+1} = A_n x_n$$

the geometry remains much more tractable.

A linear transformation can stretch, rotate, shear, contract, expand, or project a representation, but it remains within a relatively understandable geometric framework.

This gives an interesting possible interpretation of biological neural architecture:

> **Perhaps much of the apparent complexity of a nervous system comes from switching between relatively simple transformations rather than from every individual transformation being extraordinarily complicated.**

This is compatible with the basic fact that individual neurons are nonlinear while large neural populations often exhibit surprisingly structured, approximately linear population-level dynamics over particular regimes.

The nonlinearities may still be doing something important—but perhaps they do not have to perform all of the computation.

They can establish thresholds, select states, normalize signals, generate discrete events, and perform other operations while the large-scale transformation occurring inside a particular regime remains comparatively well behaved.

---

3. Even linear mappings cannot be arbitrary

This was the subtle point that emerged from the experiments.

It might initially seem that random matrices solve the problem:

> "If the mappings are linear, composition should be safe."

But it isn't.

Random matrices can themselves produce enormous geometric distortion.

After enough compositions we can get:

- exploding norms,
- collapsing dimensions,
- rapidly changing directions,
- loss of correlations,
- effective rank reduction,
- increasingly separated trajectories.

Therefore **linearity is not sufficient**.

The mappings also need some degree of **geometric compatibility**.

This leads to a stronger principle:

> **A compositional system needs transformations that preserve enough of the geometry required by subsequent transformations.**

This may be one of the more interesting lessons for biology.

---

4. A different way of thinking about neural connectivity

We often imagine a nervous system as:

neuron → neuron → neuron → neuron

and ask what function the enormous collection of synapses computes.

The concatenated mapping experiments suggest another question:

> **What kinds of transformations can a biological network safely compose for thousands or millions of computational steps?**

This shifts the focus from individual neurons toward **population geometry**.

Suppose population A represents some state x.

Population B does not necessarily need to compute some fantastically nonlinear function of x. It might primarily perform something like

$$x' \approx Ax$$

while another contextual system selects which A is appropriate.

Then we might have a sequence such as

$$x \to A_{\text{context 1}}x \to B_{\text{context 2}}A_{\text{context 1}}x \to C_{\text{context 3}}B_{\text{context 2}}A_{\text{context 1}}x$$

The resulting computation can be extremely complicated even though **each local transformation is relatively tame**.

That is close to the conceptual territory explored by context-selected linear mapping systems.

---

5. Animal brains may have enormous complexity without enormous local complexity

This may be one of the biggest biological implications.

An animal can have an astronomical number of possible behaviors without requiring an astronomical number of completely independent computations.

Instead, it could have:

a relatively small collection of useful computational transformations × many combinations of context.

For example, imagine transformations corresponding roughly to:

- orient toward stimulus,
- estimate distance,
- maintain posture,
- accelerate,
- decelerate,
- grasp,
- avoid,
- follow,
- approach,
- remember,
- compare,
- predict.

Context determines which transformations are active and how strongly they interact.

Then "hunting a mouse" does not necessarily have to be one gigantic specialized computation.

It could be a trajectory through a sequence of computational regimes.

This starts to look remarkably like a **context-selected dynamical system**.

---

6. The "islands" idea becomes especially important

A useful nervous system probably does not want its state space to be uniformly smooth.

Nor does it want it to be completely chaotic.

It needs something in between:

> **stable regions plus controlled transitions.**

One can imagine an animal's neural state space containing enormous numbers of attractor-like regions:

The points do not have to be literal fixed-point attractors. They could instead represent **temporally stable computational regimes**.

Within one region, neural activity can change continuously.

At a boundary, context changes and the system moves into another regime.

This gives us a potentially useful interpretation of things such as:

- behavioral states,
- motor programs,
- working-memory states,
- attention,
- motivational states,
- sensory contexts,
- learned associations,
- internal models.

The nervous system may therefore be organized around a very large collection of useful local regions in state space rather than operating as one globally uniform computation.

---

7. This also gives a possible explanation for recurrent structure

A purely feed-forward concatenation of mappings is simple.

Animals, however, live in time.

The actual system is therefore more naturally thought of as something like

$$x_{t+1} = A_{c_t} x_t + B_{c_t} u_t$$

where:

- x_t is internal neural state,
- u_t is sensory input,
- c_t is context.

The context can itself depend on previous state:

$$c_t = g(x_t, u_t, m_t)$$

where m_t might represent memory or a slower internal state.

This creates a **context-controlled dynamical system**.

A great deal of biological complexity can then arise without requiring every synapse to implement a complicated nonlinear computation.

The complexity comes partly from the fact that the system is continually moving through different computational regimes.

---

8. Evolution may favor composability

There is an especially interesting biological constraint here.

Evolution cannot merely search for arbitrary useful transformations. It has to produce transformations that remain **composable**.

A mutation that makes one circuit locally more powerful might actually be harmful if its output geometry becomes incompatible with the circuit that receives it.

Evolutionary pressure could therefore favor neural transformations that are:

- stable,
- approximately norm-preserving where appropriate,
- low distortion,
- redundant,
- smoothly parameterized,
- compatible with neighboring representations,
- robust to noise,

- reusable under different contexts.

This could help explain why biological systems contain so much apparent redundancy and structured connectivity.

What looks wasteful from the perspective of "compute the function as efficiently as possible" might be valuable from another perspective:

> **Keep the geometry usable by the next thousand computations.**

---

9. The "not too random" observation may be particularly important

There are really two separate sources of apparent complexity:

Combinatorial complexity

> Which mapping gets selected?

Geometric complexity

> What does that mapping actually do to the representation?

Context selection can give enormous combinatorial complexity without requiring enormous geometric complexity.

That is a powerful combination.

A system can have

$$\text{many possible behaviors}$$

while maintaining

$$\text{relatively simple local transformations.}$$

This may be a very good engineering strategy—and potentially a very good biological strategy.

---

10. A possible biological hypothesis

The overall lesson can therefore be stated as follows:

> **A large biological nervous system may achieve computational complexity less by composing arbitrarily complex transformations and more by maintaining a structured collection of relatively simple, geometrically compatible transformations whose selection and modulation are controlled by context.**

And the particularly interesting part is:

> **Context may serve not merely as information about the world, but as a mechanism for maintaining the nervous system inside a useful computational regime.**

This connects the three observations closely:

Meaningful context
→ creates stable computational regimes.

Linear or near-linear local transformations
→ makes those regimes composable.

Geometrically compatible mappings
→ prevents composition from degenerating into high-dimensional chaos.

Conclusion

The resulting picture of an animal nervous system is therefore somewhat different from the conventional picture of an enormous arbitrary nonlinear function.

It may be more useful to think of it as an:

> **enormous, context-controlled geometric dynamical system.**

Its computational power comes not only from the complexity of individual transformations, but from the enormous number of ways relatively simple transformations can be **selected, composed, reused, and placed into different contexts**.

The nervous system then has two related problems to solve.

First, it must create **useful islands of stability** in its state space.

Second, it must make the transformations within and between those islands sufficiently compatible that long chains of computation do not simply disperse into geometric chaos.

This is only a hypothesis, but it suggests an interesting direction for neuroscience: rather than asking only *what function does this neural circuit compute?*, we might also ask:

> **What geometric transformation does this circuit provide, under what context does it operate, and how does that transformation remain compatible with the transformations that follow it?**

That may be a useful way of looking for general principles governing large biological neural systems.